

# New Results on the Fault-Tolerant Facility Placement Problem

Li Yan and Marek Chrobak  
 {lyan,marek}@cs.ucr.edu  
 Computer Science, U of California Riverside  
 Riverside, CA 92521

August 30, 2011

## 1 Introduction

In the *Fault-Tolerant Facility Placement* problem (FTFP), we are given a set  $\mathcal{F}$  of *sites* at which facilities can be built, and a set  $\mathcal{C}$  of *clients* with some demands that need to be satisfied by different facilities. A client  $j$  has demand  $r_j$ . Building one facility at a site  $i$  incurs a cost  $f_i$ , and connecting one unit of demand from client  $j$  to a facility at site  $i \in \mathcal{F}$  costs  $d_{ij}$ . The distances  $d_{ij}$  form a metric, that is, they are symmetric and satisfy the triangle inequality. In a feasible solution, some number of facilities, possibly zero, are opened at each site  $i$ , and demands from each client are connected to those open facilities, with the constraint that demands from the same client have to be connected to different facilities. Note that facilities at the same site are considered different.

It is easy to see that if all  $r_j = 1$ , then FTFP reduces to the classic uncapacitated facility location problem (UFL). If we add a constraint that each site can have at most one facility built, then we get the Fault-Tolerant Facility Location problem (FTFL). One implication of the one site per facility restriction in FTFL is that  $\max_{j \in \mathcal{C}} r_j \leq |\mathcal{F}|$ , while in FTFP  $r_j$  can be much bigger than  $|\mathcal{F}|$ .

UFL has a long history and there has been great progress in designing better approximation algorithms in the past two decades. Since the publication of the first constant approximation algorithm by Shmoys, Tardos and Aardal [10], we have seen a number of techniques that are applicable in devising approximation algorithms for UFL with good approximation ratios. Using LP-rounding, Shmoys, Tardos and Aardal [10] obtained a ratio of 3.16, which was then improved by Chudak [4] to 1.736, and later by Sviridenko [11] to 1.582. Byrka [2] gave an improved LP-rounding algorithm, combined with the dual-fitting algorithm in [6], achieving a ratio of 1.5. Recently Li [9] showed that with a more refined analysis and randomizing the scaling parameter, the ratio can be improved to 1.488. This is the best known approximation result for UFL. Other techniques include Jain and Vazirani's [7] primal-dual algorithm with ratio 3, and Jain, Mahdian, Markakis, Saberi and Vazirani's [6] dual-fitting algorithm with ratio 1.61. Arya et al. [1] showed that a local search heuristic achieves a ratio of 3.

FTFL was first introduced by Jain and Vazirani [8] and they adapted their primal-dual algorithm for UFL to obtain a ratio of  $3 \ln(\max_{j \in \mathcal{C}} r_j)$ , which is logarithmic in the maximum demand. Guha, Meyerson and Monagala [5] adapted the Shmoys et al. 's [10] algorithm for UFL to obtain the first constant approximation algorithm. Swamy and Shmoys [12] improved the ratio to 2.076 using the idea of pipage rounding. Most recently, Byrka, Srinivasan and Swamy [3] improved the ratio to

1.7245 using dependent rounding and laminar clustering. All the constant approximation algorithms are based on LP-rounding.

FTFP was first introduced by Xu and Shen [13] as a generalization of UFL. They extended the dual-fitting algorithm [6] to give an approximation algorithm with a ratio claimed to be 1.861. However the algorithm runs in polynomial time only if  $\max_{j \in \mathcal{C}} r_j$  is polynomial in  $O(|\mathcal{F}| \cdot |\mathcal{C}|)$ , and the analysis of the performance guarantee seems flawed.

Our approach is to reduce FTFP to FTFL. We then use Byrka et al. 's [3] result on FTFL to achieve the ratio 1.7245.

## 2 The LP Formulation

The FTFP problem has a natural IP formulation. Let  $y_i$  represent the number of facilities built at site  $i$  and let  $x_{ij}$  represent the number of connections from client  $j$  to facilities at site  $i$ . If we relax the integral constraints, we obtain the following LP:

$$\begin{aligned}
& \text{minimize} && \sum_{i \in \mathcal{F}} f_i y_i + \sum_{i \in \mathcal{F}, j \in \mathcal{C}} d_{ij} x_{ij} && (1) \\
& \text{subject to} && y_i - x_{ij} \geq 0, && \forall i \in \mathcal{F}, j \in \mathcal{C} \\
& && \sum_{i \in \mathcal{F}} x_{ij} \geq r_j, && \forall j \in \mathcal{C} \\
& && x_{ij} \geq 0, y_i \geq 0, && \forall i \in \mathcal{F}, j \in \mathcal{C}
\end{aligned}$$

The dual program is:

$$\begin{aligned}
& \text{maximize} && \sum_{j \in \mathcal{C}} r_j \alpha_j && (2) \\
& \text{subject to} && \sum_{j \in \mathcal{C}} \beta_{ij} \leq f_i, && \forall i \in \mathcal{F} \\
& && \alpha_j - \beta_{ij} \leq d_{ij}, && \forall i \in \mathcal{F}, j \in \mathcal{C} \\
& && \alpha_j \geq 0, \beta_{ij} \geq 0, && \forall i \in \mathcal{F}, j \in \mathcal{C}
\end{aligned}$$

## 3 A 1.7245-Approximation Algorithm

In this section we give an algorithm based on a reduction to FTFL. Moreover, we show that if we have a  $\rho$ -approximation algorithm for FTFL, then we can use it to build an approximation algorithm for FTFP with the same ratio. To be more precise, we assume that  $\rho$  is the ratio of the cost of the FTFL solution to the fractional optimal solution cost of the natural LP formulation for FTFL.<sup>1</sup>

A naïve idea is to split the sites so that each site is split into  $P = \max_{j \in \mathcal{C}} r_j$  identical sites. Now we can restrict each split site to have at most one facility. Since we never need more than  $P$  facilities open at the same site in the original instance, the FTFL instance after split is equivalent to the original FTFP instance. However, since  $P$  might be large, this reduction might result in an instance with exponential size.

---

<sup>1</sup>Adding a constraint  $y_i \leq 1, \forall i \in \mathcal{F}$  to the LP (1) results in the LP for FTFL.

However, we can modify the naïve approach to obtain an FTFL instance with polynomial size, using an optimal fractional solution to the LP (1).

Let  $(\mathbf{x}^*, \mathbf{y}^*)$  be the fractional optimal solution to LP (1). Also let

$$\hat{y}_i = \max\{0, \lfloor y_i^* \rfloor - 1\}, \quad \hat{x}_{ij} = \min\{\lfloor x_{ij}^* \rfloor, \hat{y}_i\},$$

and

$$\bar{x}_{ij} = x_{ij}^* - \hat{x}_{ij}, \quad \bar{y}_i = y_i^* - \hat{y}_i.$$

Let  $\mathcal{I}$  be the original FTFL instance. We define an FTFL instance  $\mathcal{I}_1$  with the same set  $\mathcal{F}$  of sites and  $\mathcal{C}$  of clients, the same distances  $d_{ij}$  and the same facility costs  $f_i$ , except that the demands are  $\hat{r}_j = \sum_{i \in \mathcal{F}} \hat{x}_{ij}$ . Let  $S_1$  be the solution defined by  $(\hat{x}_{ij}, \hat{y}_i)$ . Clearly  $S_1$  is a feasible integral solution to the instance  $\mathcal{I}_1$ .

Another FTFL instance  $\mathcal{I}_2$  is defined using the same parameters except that the demands are  $\bar{r}_j = r_j - \hat{r}_j = \sum_{i \in \mathcal{F}} \bar{x}_{ij}$ . We claim that  $(\bar{x}_{ij}, \bar{y}_i)$  is a feasible fractional solution to the instance  $\mathcal{I}_2$ <sup>2</sup>. To ensure feasibility, we only need to show  $\bar{x}_{ij} \leq \bar{y}_i$  for all  $i \in \mathcal{F}, j \in \mathcal{C}$ . The proof proceeds by considering two cases. Case 1 is that  $y_i^* < 1$ . Then we have  $\bar{x}_{ij} = x_{ij}^* \leq y_i^* = \bar{y}_i$  as needed. Case 2 is  $y_i^* \geq 1$ . From the definition we have  $\bar{y}_i \geq 1$ . We further consider two subcases: if  $\hat{x}_{ij} = \hat{y}_i$ , then  $\bar{x}_{ij} = x_{ij}^* - \hat{x}_{ij} \leq y_i^* - \hat{y}_i \leq \bar{y}_i$  and we are done. The other subcase is  $\hat{x}_{ij} < \hat{y}_i$ . It follows that  $0 \leq \bar{x}_{ij} < 1$ . Again we have  $\bar{x}_{ij} < 1 \leq \bar{y}_i$ . The above takes care of all cases, so we have shown the feasibility of the fractional solution  $(\bar{x}_{ij}, \bar{y}_i)$  to the instance  $\mathcal{I}_2$ .

One more observation about the instance  $\mathcal{I}_2$  is that  $\bar{r}_j \leq 2n$  where  $n = |\mathcal{F}|$ , since  $\bar{r}_j = \sum_{i \in \mathcal{F}} \bar{x}_{ij}$  and for every  $i \in \mathcal{F}, j \in \mathcal{C}$ , we have  $0 \leq \bar{x}_{ij} < 2$ . Therefore we have  $\bar{P} = \max_{j \in \mathcal{C}} \bar{r}_j \leq 2n$ . Now that we have the demands being polynomial w.r.t. the input size of the original FTFL instance, we can split each site into  $\bar{P}$  new sites and treat the instance as an FTFL instance. Solving the FTFL instance with a  $\rho$ -approximation algorithm, we have an integral solution  $S_2$  with cost no more than  $\rho \cdot \text{LP}^*(\mathcal{I}_2)$ , where  $\text{LP}^*(\mathcal{I}_2)$  is the cost of a fractional optimal solution to the instance  $\mathcal{I}_2$ .

Combining  $S_1$  and  $S_2$ , we have a feasible solution to the original FTFL instance. We now argue that the solution is within a factor of  $\rho$  from the cost of an optimal fractional solution to the LP (1).

First we observe that  $\text{cost}(S_1) = \sum_{i \in \mathcal{F}} f_i \hat{y}_i + \sum_{i \in \mathcal{F}, j \in \mathcal{C}} d_{ij} \hat{x}_{ij}$ . Secondly, since  $(\bar{x}_{ij}, \bar{y}_i)$  is a feasible fractional solution to the instance  $\mathcal{I}_2$ , we have  $\sum_{i \in \mathcal{F}} f_i \bar{y}_i + \sum_{i \in \mathcal{F}, j \in \mathcal{C}} d_{ij} \bar{x}_{ij} \geq \text{LP}^*(\mathcal{I}_2)$ . Our solution has total cost

$$\begin{aligned} \text{cost}(S_1) + \text{cost}(S_2) &\leq \sum_{i \in \mathcal{F}} f_i \hat{y}_i + \sum_{i \in \mathcal{F}, j \in \mathcal{C}} d_{ij} \hat{x}_{ij} + \rho \cdot \text{LP}^*(\mathcal{I}_2) \\ &\leq \sum_{i \in \mathcal{F}} f_i \hat{y}_i + \sum_{i \in \mathcal{F}, j \in \mathcal{C}} d_{ij} \hat{x}_{ij} + \rho \left( \sum_{i \in \mathcal{F}} f_i \bar{y}_i + \sum_{i \in \mathcal{F}, j \in \mathcal{C}} d_{ij} \bar{x}_{ij} \right) \\ &\leq \rho \left( \sum_{i \in \mathcal{F}} f_i \hat{y}_i + \sum_{i \in \mathcal{F}, j \in \mathcal{C}} d_{ij} \hat{x}_{ij} + \sum_{i \in \mathcal{F}} f_i \bar{y}_i + \sum_{i \in \mathcal{F}, j \in \mathcal{C}} d_{ij} \bar{x}_{ij} \right) \\ &= \rho \left( \sum_{i \in \mathcal{F}} f_i y_i^* + \sum_{i \in \mathcal{F}, j \in \mathcal{C}} d_{ij} x_{ij}^* \right) = \rho \cdot \text{LP}^*(\mathcal{I}) \\ &\leq \rho \cdot \text{OPT}(\mathcal{I}) \end{aligned}$$

---

<sup>2</sup>Note that if we take  $\hat{x}_{ij} = \lfloor x_{ij}^* \rfloor, \hat{y}_i = \lfloor y_i^* \rfloor$  and define  $\bar{x}_{ij}, \bar{y}_i, \bar{r}_j$  in a similar way, then  $(\bar{x}_{ij}, \bar{y}_i)$  may not be feasible to the instance with demands  $\{\bar{r}_j\}$  because we might have  $\bar{x}_{ij} > \bar{y}_i$  if  $0 < x_{ij}^* < 1$  and  $y_i^* = 1$ .

The currently best known approximation algorithm for FTFL achieves a ratio of 1.7245 [3], so our algorithm is a 1.7245-approximation algorithm for FTFP.

## 4 Approximation for Large Demands

If we seek an approximation algorithm with ratio as a function of  $R = \min_{j \in \mathcal{C}} r_j$ , then we show that we can achieve a ratio of  $1 + O(n/R)$  where  $n = |\mathcal{F}|$ . The hidden constant in the big-O is the same as the best approximation ratio for FTFL, which is known to be at most 1.7245 [3].<sup>3</sup>

Our algorithm uses the 1.7245-approximation algorithm for FTFL as a subroutine. First we solve the LP (1) and obtain  $(x_{ij}^*, y_i^*)$  as the optimal fractional solution. Now we round down the fractional values as before and the rounded down values  $\hat{x}_{ij} = \lfloor x_{ij}^* \rfloor, \hat{y}_i = \lfloor y_i^* \rfloor$  form part of the final solution, denoted as  $S_1$ . Now each client  $j$  has a reduced demand  $\bar{r}_j = r_j - \sum_{i \in \mathcal{F}} \hat{x}_{ij} \leq n - 1$ , where  $n = |\mathcal{F}|$ . Therefore we can solve this instance with reduced demands, denoted as  $\mathcal{I}_2$ , by solving an equivalent FTFL instance: We replicate each site with  $n - 1$  copies since no site needs to have more than  $n - 1$  facilities open. Now we can restrict each duplicated site open at most one facility and this gives us an FTFL instance with polynomial size to work with. We then use the 1.7245-approximation algorithm to solve this FTFL instance to obtain an integral solution  $S_2$  with  $\text{cost}(S_2) \leq \rho \cdot \text{LP}^*(\mathcal{I}_2)$  with  $\rho = 1.7245$ . We now show that  $\text{LP}^*(\mathcal{I}_2) \leq \frac{n}{R} \text{LP}^*(\mathcal{I}) \leq \frac{n}{R} \text{OPT}(\mathcal{I})$ , where  $\mathcal{I}$  is the input FTFP instance. To see this, consider three FTFP instances:  $\mathcal{I}_1$  is the instance with all  $r_j = R$ , and  $\mathcal{I}_2$  is the instance with all  $r_j = \bar{r}_j$ , and  $\mathcal{I}_3$  is the instance with all demands  $r_j = n - 1$ . We have  $\text{LP}^*(\mathcal{I}_3)/\text{LP}^*(\mathcal{I}_1) = (n - 1)/R$  since any fractional solution with cost  $Z$  of an FTFP instance with uniform demands  $s$  can be scaled down by  $s$  to obtain a feasible solution for the same instance with demands set to 1 and the cost being  $Z/s$ . We also have  $\text{LP}^*(\mathcal{I}_3) \geq \text{LP}^*(\mathcal{I}_2)$  since every demand in  $\mathcal{I}_3$  is at least as large as that in  $\mathcal{I}_2$ . Similarly  $\text{LP}^*(\mathcal{I}) \geq \text{LP}^*(\mathcal{I}_1)$ . Therefore we have

$$\frac{\text{LP}^*(\mathcal{I}_2)}{\text{LP}^*(\mathcal{I})} \leq \frac{\text{LP}^*(\mathcal{I}_3)}{\text{LP}^*(\mathcal{I}_1)} = \frac{n - 1}{R} \leq \frac{n}{R}.$$

By combining  $S_1$  and  $S_2$  we have a feasible solution to the original FTFP instance. Now we bound the total cost. First notice that  $\text{cost}(S_1) \leq \text{LP}^*(\mathcal{I}) \leq \text{OPT}(\mathcal{I})$ . By the previous argument, we have  $\text{cost}(S_2) \leq \rho \frac{n}{R} \text{OPT}(\mathcal{I})$ . Therefore, the total cost of our solution is

$$\text{cost}(S_1) + \text{cost}(S_2) \leq (1 + \rho \frac{n}{R}) \text{OPT}(\mathcal{I}),$$

which is  $(1 + O(\frac{n}{R})) \text{OPT}(\mathcal{I})$ . The approximation ratio approaches 1 when  $R$  becomes much larger than  $n$ . This is perhaps not too surprising as when all the demands are large, the fractional optimal solution to the LP (1) is close to integral relative to the demands, so their costs are close.

## References

- [1] Vijay Arya, Naveen Garg, Rohit Khandekar, Adam Meyerson, Kamesh Munagala, and Vinayaka Pandit. Local search heuristic for k-median and facility location problems. In *Proc. 33rd Symp. Theory of Computing (STOC)*, pages 21–29, New York, NY, USA, 2001. ACM.

---

<sup>3</sup>The hidden constant can be improved to be 1.488 by solving a sequence of UFL instances.

- [2] Jaroslaw Byrka and Karen Aardal. An optimal bifactor approximation algorithm for the metric uncapacitated facility location problem. *SIAM J. Comput.*, 39(6):2212–2231, 2010.
- [3] Jaroslaw Byrka, Aravind Srinivasan, and Chaitanya Swamy. Fault-tolerant facility location, a randomized dependent LP-rounding algorithm. In *Proc. 14th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 244–257, 2010.
- [4] Fabián Chudak and David Shmoys. Improved approximation algorithms for the uncapacitated facility location problem. *SIAM J. Comput.*, 33(1):1–25, January 2004.
- [5] Sudipto Guha, Adam Meyerson, and Kamesh Munagala. Improved algorithms for fault tolerant facility location. In *Proc. 12th Symp. on Discrete Algorithms (SODA)*, pages 636–641, 2001.
- [6] Kamal Jain, Mohammad Mahdian, Evangelos Markakis, Amin Saberi, and Vijay Vazirani. Greedy facility location algorithms analyzed using dual fitting with factor-revealing LP. *J. ACM*, 50(6):795–824, November 2003.
- [7] Kamal Jain and Vijay Vazirani. Approximation algorithms for metric facility location and k-median problems using the primal-dual schema and lagrangian relaxation. *J. ACM*, 48(2):274–296, 2001.
- [8] Kamal Jain and Vijay V. Vazirani. An approximation algorithm for the fault tolerant metric facility location problem. *Algorithmica*, 38(3):433–439, 2003.
- [9] Shi Li. A 1.488 approximation algorithm for the uncapacitated facility location problem. In Luca Aceto, Monika Henzinger, and Jiri Sgall, editors, *Proc. 38rd International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 6756, pages 77–88. Springer, 2011.
- [10] David Shmoys, Éva Tardos, and Karen Aardal. Approximation algorithms for facility location problems (extended abstract). In *Proc. 29th Symp. Theory of Computing (STOC)*, pages 265–274, New York, NY, USA, 1997. ACM.
- [11] Maxim Sviridenko. An improved approximation algorithm for the metric uncapacitated facility location problem. In *Proc. 9th Conference on Integer Programming and Combinatorial Optimization (IPCO)*, pages 240–257, London, UK, 2002. Springer-Verlag.
- [12] Chaitanya Swamy and David Shmoys. Fault-tolerant facility location. *ACM Trans. Algorithms*, 4(4):1–27, August 2008.
- [13] Shihong Xu and Hong Shen. The fault-tolerant facility allocation problem. In *Proc. 20th International Symp. on Algorithms and Computation (ISAAC)*, pages 689–698, Berlin, Heidelberg, 2009. Springer-Verlag.